

PATENT
Attorney's Docket No. **0023-0200**

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of:)	MAIL STOP APPEAL BRIEF - PATENTS
)	
Changming LIU et al.)	Group Art Unit: 2143
)	
Application No.: 09/658,424)	Examiner: D. England
)	
Filed: September 8, 2000)	
)	
For: GUARANTEED BANDWIDTH)	
SHARING IN A TRAFFIC)	
SHAPING SYSTEM)	

United States Patent & Trademark Office
Customer Service Window, Mail Stop Appeal Brief - Patents
Randolph Building
401 Dulany Street
Alexandria, Virginia 22314

APPEAL BRIEF

This Appeal Brief is submitted in response to the final Office Action, dated January 26, 2007, and in support of the Notice of Appeal, filed March 9, 2007.

I. **REAL PARTY IN INTEREST**

The real party in interest in this appeal is Juniper Networks, Inc.

II. RELATED APPEALS, INTERFERENCES, AND JUDICIAL PROCEEDINGS

Appellants are unaware of any related appeals, interferences or judicial proceedings.

III. STATUS OF CLAIMS

Claims 1-22 are pending in this application. Claims 1-22 were rejected in the final Office Action, dated January 26, 2007, and are the subject of the present appeal. These claims are reproduced in the Claim Appendix of this Appeal Brief.

IV. STATUS OF AMENDMENTS

No amendment was filed subsequent to the final Office Action, dated January 26, 2007.

V. SUMMARY OF CLAIMED SUBJECT MATTER

In the paragraphs that follow, a concise explanation of the independent claims, each dependent claim argued separately, and the claims reciting means-plus-function or step-plus-function language that are involved in this appeal will be provided by referring, in parenthesis, to examples of where support can be found in the specification and drawings.

Claim 1 is directed to a method for allocating bandwidth in a network appliance where the network appliance includes a plurality of guaranteed bandwidth buckets used to evaluate when to pass traffic through the network appliance (e.g., 100, Fig. 1; pg. 5,

lines 13-17; pg. 6, lines 1-3), the method includes providing a shared bandwidth bucket associated with each of the plurality of the guaranteed bandwidth buckets (e.g., 130c, 130b, Fig. 1; pg. 6, lines 10-30 to pg. 7, lines 1-3); allocating bandwidth to the shared bandwidth bucket based on the underutilization of bandwidth in any one of the plurality of guaranteed bandwidth buckets (e.g., pg. 7, lines 3-7); determining whether bandwidth in one of the plurality of guaranteed bandwidth buckets is sufficient to allow traffic to pass immediately through the network appliance (e.g., pg. 7, lines 10-12); and transferring bandwidth from the shared bandwidth bucket to one of the plurality of guaranteed bandwidth buckets when it is determined that bandwidth in one of the plurality of guaranteed bandwidth buckets is not sufficient to allow traffic to pass immediately through the network appliance (e.g., pg. 7, lines 14-16).

Claim 14 is directed to a method for allocating bandwidth in a network appliance including defining a guaranteed bandwidth allocation for a first policy for passing traffic through the network appliance (e.g., 100, Fig. 1; pg. 5, lines 13-17; pg. 6, lines 1-3) including using a first bucket to allocate the guaranteed bandwidth (e.g., 130a, Fig. 1; pg. 6, lines 10-15); defining a guaranteed bandwidth allocation for a second policy for passing traffic through the network appliance including using a second bucket to allocate the guaranteed bandwidth (e.g., 130b, Fig. 1; pg. 6, lines 21-30); sharing excess bandwidth developed from the underutilization of the guaranteed bandwidth allocated to the first and second buckets including providing a shared bandwidth bucket associated with the first and second buckets (e.g., 130c, Fig. 1; pg. 7, lines 1-7); and borrowing bandwidth from the shared bandwidth bucket by one of the first and second buckets when

the respective bucket has insufficient bandwidth to allow traffic to pass immediately through the network appliance (e.g., pg. 7, lines 10-16).

Claim 15 is directed to an apparatus for allocating bandwidth in a network appliance (e.g., 100, Fig. 1; pg. 5, lines 13-17; pg. 6, lines 1-3) where the network appliance includes a plurality of guaranteed bandwidth buckets used to evaluate when to pass traffic through the network appliance (e.g., 130b, Fig. 1; pg. 6, lines 21-30), the apparatus includes a shared bandwidth bucket (e.g., 130c, Fig. 1; pg. 7, lines 1-7) associated with a plurality of the guaranteed bandwidth buckets; means for allocating bandwidth to the shared bandwidth bucket based on the underutilization of bandwidth in the plurality of guaranteed bandwidth buckets (e.g., pg. 7, lines 10-16); and a scheduler (e.g., 108, Fig. 1; pg. 5, lines 20-24) operable to evaluate a packet to determine if a traffic shaping policy should be applied to a given packet (e.g., 206, Fig. 2; pg. 7, lines 25-26), evaluate a guaranteed bandwidth bucket associated with an identified traffic shaping policy (e.g., 210, Fig. 2; pg. 8, lines 5-10), determine when the guaranteed bandwidth bucket associated with an identified traffic shaping policy has insufficient capacity to support a transfer of the packet through the network (e.g., 210, Fig. 2; pg. 8, lines 5-10), and borrow bandwidth from the shared bandwidth bucket by a respective guaranteed bandwidth bucket to allow traffic to pass immediately through the network appliance (e.g., 214, Fig. 2; pg. 8, lines 15-18).

Claim 16 is directed to a network device (e.g., 100, Fig. 1; pg. 5, lines 13-17; pg. 6, lines 1-3) including a first bucket configured to receive tokens at a first information rate (e.g., 130a, Fig. 1; pg. 6, lines 10-15); a second bucket configured to receive tokens at a second information rate (e.g., 130b, Fig. 1; pg. 6, lines 21-30); a third bucket

configured to receive extra tokens from the second bucket (e.g., 130c, Fig. 1; pg. 7, lines 1-7); and a scheduler (e.g., 108, Fig. 1; pg. 5, lines 20-24) configured to: determine if a size of traffic received at the network device exceeds a number of tokens stored in the first bucket (e.g., 208, Fig. 2; pg. 8, lines 3-7), determine, when the size of the traffic does not exceed the number of tokens stored in the first bucket, if a size of the traffic exceeds a number of tokens stored in the second bucket (e.g., 210, Fig. 2; pg. 8, lines 5-10), and transfer, when the size of the traffic exceeds the number of tokens stored in the second bucket, an appropriate number of tokens from the third bucket to the second bucket so that the second bucket includes a number of tokens that equals or exceeds the size of the traffic (e.g., 214, Fig. 2; pg. 8, lines 15-18).

Claim 20 is directed to a method including receiving traffic; determining if a policy is to be applied to the traffic (e.g., 206, Fig. 2; pg. 7, lines 25-26); determining, when a policy is to be applied to the traffic, if a size of the traffic exceeds a number of tokens in a first bucket, the first bucket being associated with the policy (e.g., 208, Fig. 2; pg. 8, lines 3-7); determining, when the size of the traffic does not exceed the number of tokens in the first bucket, if the size of the traffic exceeds the number of tokens in a second bucket (e.g., 210, Fig. 2; pg. 8, lines 5-10); determining, when the size of the traffic exceeds the number of tokens in the second bucket, if a third bucket includes an appropriate number of tokens that, when added to the number of tokens in the second bucket, would equal or exceed the size of the traffic (e.g., 212, Fig. 2; pg. 8, lines 11-13); transferring the appropriate number of tokens from the third bucket to the second bucket when the third bucket includes the appropriate number of tokens; and forwarding the traffic after the transferring (e.g., 214, Fig. 2; pg. 8, lines 15-18).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

A. Claims 1, 5, 6 and 14 stand rejected under 35 U.S.C. § 102(e) as being anticipated by IVERSON et al. (U.S. Patent No. 6,052,379).

B. Claims 2, 3, 7-11, 13 and 15-22 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over IVERSON et al. in view of HO (U.S. Patent No. 6,862,270).

C. Claim 4 stands rejected under 35 U.S.C. § 103(a) as being unpatentable over IVERSON et al. in view of Applicants' admitted prior art.

D. Claim 12 stands rejected under 35 U.S.C. § 103(a) as allegedly unpatentable over IVERSON et al. in view of CHIRUVOLU (U.S. Patent No. 6,839,321).

VII. ARGUMENTS**A. The rejection under 35 U.S.C. § 102(e) based on IVERSON et al. (U.S. Patent No. 6,052,379) should be reversed.**

The initial burden of establishing a *prima facie* basis to deny patentability to a claimed invention always rests upon the Examiner. In re Oetiker, 977 F.2d 1443, 24 U.S.P.Q.2d 1443 (Fed. Cir. 1992). A proper rejection under 35 U.S.C. § 102 requires that a single reference teach every aspect of the claimed invention. Any feature not directly taught must be inherently present. Verdegaal Bros. v. Union Oil Co. of California, 814 F.2d 628, 2 USPQ2d 1051 (Fed. Cir. 1987).

1. Claims 1, 5, and 6.

Independent claim 1 is directed to a method for allocating bandwidth in a network appliance where the network appliance includes a plurality of guaranteed bandwidth buckets used to evaluate when to pass traffic through the network appliance, the method includes providing a shared bandwidth bucket associated with each of the plurality of the guaranteed bandwidth buckets; allocating bandwidth to the shared bandwidth bucket based on the underutilization of bandwidth in any one of the plurality of guaranteed bandwidth buckets; determining whether bandwidth in one of the plurality of guaranteed bandwidth buckets is sufficient to allow traffic to pass immediately through the network appliance; and transferring bandwidth from the shared bandwidth bucket to one of the plurality of guaranteed bandwidth buckets when it is determined that bandwidth in one of the plurality of guaranteed bandwidth buckets is not sufficient to allow traffic to pass immediately through the network appliance. IVERSON et al. does not disclose or suggest this combination of features.

For example, IVERSON et al. does not disclose or suggest providing a shared bandwidth bucket associated with each of a plurality of the guaranteed bandwidth buckets. The Examiner relies on the Abstract, Fig. 10, and col. 17, line 56 to col. 18, line 19 of IVERSON et al. for allegedly disclosing this feature (final Office Action -- pg. 2). Appellants respectfully disagree with the Examiner's interpretation of IVERSON et al.

The Abstract of IVERSON et al. discloses:

A priority scheme is based on an amount of preallocated bandwidth unused by channel unit ports. A first water level in a first bucket is associated with an amount of allotted bandwidth unused by the channel unit and a second water level in a second bucket is associated with an amount of unused allotted bandwidth exceeding an overflow level of the first bucket. A priority value is derived from the first water level when the first water level is above zero. The priority value is derived from the second water level when the first water level is below or equal to

zero. In another aspect of the invention, the high priority value is determined by tracking a percentage utilization of allocated bandwidth for a predetermined number of time increments comprising a measurement time period.

This section of IVERSON et al. discloses a first allotted bandwidth bucket and a second overflow bucket. Depending on whether a water level in the first bucket is above or below zero, a priority value is derived from the level of either the first bucket or the second bucket. This section of IVERSON et al. does not disclose or suggest providing a shared bandwidth bucket associated with a **plurality** of guaranteed bandwidth buckets, as recited in claim 1.

Col. 17, line 56 – col. 18, line 19 of IVERSON et al. discloses:

If the BpCSum is positive, the port was requesting bandwidth at a rate below the CIR+B_c for at least the last measurement interval. If the BpCSum is zero, port bandwidth requests have been substantially equal to the CIR+B_c for the port. If the water level in CSum is negative (below the midpoint), the rate that the port has been using bandwidth is above CIR+B_c. If the port has accumulated any excess bandwidth credit by transmitting below CIR for some amount of time, this bandwidth credit will be used if the water level in the first bucket goes below zero.

BpEsum is the water level value in the second bucket 404 and represents the current accumulated value of unused bandwidth in excess of CIR+B_c (i.e. past overflows from the first bucket 402). The Esum bucket 404 represents a cache of excess bandwidth that the user 62 can save up to be used for longer periods of high transmission demand.

Every measurement interval the quantum of bits 400 are added to the first bucket 402. Any overflow of bandwidth above the limit of the first bucket 402 is added to the Esum bucket 404.

Both buckets are "leaky" in that the amount of traffic transmitted in the past measurement interval leaks out of the appropriate bucket based on the previous priority level. The current water level of each bucket is then the result of adding in the Committed Information Rate (CIR) bit quantum for the last measurement interval and subtracting the amount of outgoing traffic 409 actually transmitted in the last measurement interval, T1Out. The water level of bucket 402 determines a priority value in a high priority

band 403. The water level of bucket 404 determines a priority value in a low priority band 405.

This section of IVERSON et al. discloses a leaky bucket priority scheme, wherein excess bandwidth credits for a first committed bandwidth bucket 402 are added to a second excess bandwidth bucket 404. The excess bandwidth stored in bucket 404 is then used when the level of the first bucket 402 drops below zero (a midpoint in the bucket). This section of IVERSON et al. does not disclose or suggest providing a shared bandwidth bucket associated with a **plurality** of guaranteed bandwidth buckets, as recited in claim

1. Even assuming *arguendo* that IVERSON et al. disclose a shared bandwidth bucket (e.g., second bucket 404) associated with a single guaranteed bandwidth bucket (e.g., first bucket 402), a point that Appellants do not concede, this association is clearly a one-to-one association, resulting in bandwidth overages from bucket 402 being applied to bucket 404 for subsequent use when the level of bucket 402 drops below zero. Contrary to this disclosure, claim 1 recites a shared bandwidth bucket being associated with a **plurality of guaranteed bandwidth buckets**. By associating multiple guaranteed bandwidth buckets with a shared bandwidth bucket, traffic resources may be more optimally distributed.

Clearly, IVERSON et al. fails to disclose each and every element of claim 1, as required under 35 U.S.C. § 102.

In responding to Appellants arguments relating to claim 1, the Examiner indicates that the “first bucket” in Iverson is the bucket 404, while the plurality of guaranteed bandwidth buckets may be interpreted as CIR and bucket 402. (final Office Action -- pg. 11). Following through on this rationale, equating the system of IVERSON et al. to the method of claim 1, IVERSON et al. must disclose, either explicitly or inherently,

allocating bandwidth to the bucket 404 based on the underutilization of bandwidth in the CIR 400 and the first bucket 402; and transferring bandwidth developed from the underutilization of the guaranteed bandwidth allocated to CIR 400 and first bucket 402 including borrowing bandwidth from the second bucket 404 by a respective guaranteed bandwidth bucket (i.e., CIR 400 and first bucket 402) to allow traffic to pass immediately through the network appliance.

Clearly, IVERSON et al. does not disclose or even remotely suggest allocating bandwidth to the second bucket 404 based on the underutilization of bandwidth in CIR 400. On the contrary, all bandwidth delivered by CIR 400 is ‘used’ in terms of its allocation to a port. This is the very nature of the committed information rate bit quantum. At col. 17, lines 40-42, IVERSON et al. discloses “[a]t the end of every evaluation interval *the Committed Information Rate (CIR) quantum is emptied* into a the (sic) CSum bucket 402 and/or the ESum bucket 404.” (emphasis added). In this manner, it remains clear that allocation of bandwidth to the second bucket 404 is not based on the underutilization of bandwidth in CIR 404, as suggested by the Examiner. As described above, second bucket 404 is associated directly with first bucket 402 to maintain excess bandwidth allocated to, but not used by, first bucket 402.

Furthermore, in responding to Appellants prior remarks, the Examiner indicated that the mere duplication of essential working parts of a device involves only routine skill in the art (citing St. Regis Paper Co. v. Bemis Co., 193 USPQ 8; Office Action date December 20, 2005 – pg. 11). It should be initially noted, that the cited “rule” relates to findings of obviousness rather than anticipation, since rejections under 35 U.S.C. §102 must disclose each and every feature of a claimed invention. Accordingly, failing to

disclose a claimed feature, even one alleged to be duplicative, prevents application of §102.

Additionally, the Examiner does not compare the facts in St. Regis Paper Co. with those in the present case and explain why, based upon this comparison, the legal conclusion in the present case should be the same as that in St. Regis Paper Co. Instead, the examiner relies upon St. Regis Paper Co. as establishing a per se rule that duplication of parts involves only routine skill in the art. As stated by the Federal Circuit in In re Ochiai, 71 F.3d 1565, 1572, 37 USPQ2d 1127, 1133 (Fed. Cir. 1995), “reliance on per se rules of obviousness is legally incorrect and must cease.”

Claims 5 and 6 depend from claim 1. Therefore, these claims are patentable over IVERSON et al. for at least the reasons given above with respect to claim 1.

For at least the foregoing reasons, Appellants submit that the rejection of claims 1, 5, and 6 under 35 U.S.C. § 102(e) based on IVERSON et al. are improper. Accordingly, Appellants request that the rejection be reversed.

2. Claim 14.

Independent claim 14 is directed to a method for allocating bandwidth in a network appliance including defining a guaranteed bandwidth allocation for a first policy for passing traffic through the network appliance including using a first bucket to allocate the guaranteed bandwidth; defining a guaranteed bandwidth allocation for a second policy for passing traffic through the network appliance including using a second bucket to allocate the guaranteed bandwidth; sharing excess bandwidth developed from the underutilization of the guaranteed bandwidth allocated to the first and second buckets including providing a shared bandwidth bucket associated with the first and second

buckets; and borrowing bandwidth from the shared bandwidth bucket by one of the first and second buckets when the respective bucket has insufficient bandwidth to allow traffic to pass immediately through the network appliance. IVERSON et al. does not disclose or suggest this combination of features.

For example, IVERSON et al. does not disclose or suggest providing a shared bandwidth bucket associated with the first and second buckets. Although not explicitly stated in the final Office Action, the Examiner appears to again rely on the Abstract, Fig. 10, and col. 17, line 56 to col. 18, line 19 of IVERSON et al. for allegedly disclosing this feature (final Office Action -- pg. 3 ("Claim 14 is rejected for similar reasons as stated above"). Appellants respectfully disagree with the Examiner's interpretation of IVERSON et al.

The Abstract of IVERSON et al. discloses:

A priority scheme is based on an amount of preallocated bandwidth unused by channel unit ports. A first water level in a first bucket is associated with an amount of allotted bandwidth unused by the channel unit and a second water level in a second bucket is associated with an amount of unused allotted bandwidth exceeding an overflow level of the first bucket. A priority value is derived from the first water level when the first water level is above zero. The priority value is derived from the second water level when the first water level is below or equal to zero. In another aspect of the invention, the high priority value is determined by tracking a percentage utilization of allocated bandwidth for a predetermined number of time increments comprising a measurement time period.

This section of IVERSON et al. discloses a first allotted bandwidth bucket and a second overflow bucket. Depending on whether a water level in the first bucket is above or below zero, a priority value is derived from the level of either the first bucket or the second bucket. This section of IVERSON et al. does not disclose or suggest providing a

shared bandwidth bucket associated with the **first and second buckets**, as recited in claim 14.

At col. 17, line 56 – col. 18, line 19 of IVERSON et al. discloses:

If the BpCSum is positive, the port was requesting bandwidth at a rate below the CIR+B_c for at least the last measurement interval. If the BpCSum is zero, port bandwidth requests have been substantially equal to the CIR+B_c for the port. If the water level in CSum is negative (below the midpoint), the rate that the port has been using bandwidth is above CIR+B_c. If the port has accumulated any excess bandwidth credit by transmitting below CIR for some amount of time, this bandwidth credit will be used if the water level in the first bucket goes below zero.

BpEsum is the water level value in the second bucket 404 and represents the current accumulated value of unused bandwidth in excess of CIR+B_c (i.e. past overflows from the first bucket 402). The ESum bucket 404 represents a cache of excess bandwidth that the user 62 can save up to be used for longer periods of high transmission demand.

Every measurement interval the quantum of bits 400 are added to the first bucket 402. Any overflow of bandwidth above the limit of the first bucket 402 is added to the ESum bucket 404.

Both buckets are "leaky" in that the amount of traffic transmitted in the past measurement interval leaks out of the appropriate bucket based on the previous priority level. The current water level of each bucket is then the result of adding in the Committed Information Rate (CIR) bit quantum for the last measurement interval and subtracting the amount of outgoing traffic 409 actually transmitted in the last measurement interval, T1Out. The water level of bucket 402 determines a priority value in a high priority band 403. The water level of bucket 404 determines a priority value in a low priority band 405.

This section of IVERSON et al. discloses a leaky bucket priority scheme, wherein excess bandwidth credits for a first committed bandwidth bucket 402 are added to a second excess bandwidth bucket 404. The excess bandwidth stored in bucket 404 is then used when the level of the first bucket 402 drops below zero (a midpoint in the bucket). This section of IVERSON et al. does not disclose or suggest providing a shared bandwidth

bucket associated with the first and second buckets, as recited in claim 14. Even assuming *arguendo* that IVERSON et al. disclose a shared bandwidth bucket (e.g., second bucket 404) associated with a single guaranteed bandwidth bucket (e.g., first bucket 402), a point that Appellants do not concede, this association is clearly a one-to-one association, resulting in bandwidth overages from bucket 402 being applied to bucket 404 for subsequent use when the level of bucket 402 drops below zero. Contrary to this disclosure, claim 14 recites providing a shared bandwidth bucket associated with **the first and second buckets**. By associating multiple guaranteed bandwidth buckets with a shared bandwidth bucket, traffic resources may be more optimally distributed. For at least this reason, IVERSON et al. fails to disclose each and every element of claim 14, as required under 35 U.S.C. § 102.

Furthermore, IVERSON et al. does not disclose or suggest defining a guaranteed bandwidth allocation for a first policy for passing traffic through the network appliance including using a first bucket to allocate the guaranteed bandwidth and defining a guaranteed bandwidth allocation for a second policy for passing traffic through the network appliance including using a second bucket to allocate the guaranteed bandwidth and borrowing bandwidth from the shared bandwidth bucket by one of the first and second buckets when the respective bucket has insufficient bandwidth to allow traffic to pass immediately through the network appliance, as recited in claim 14.

As in prior Office Actions, the Examiner continues to **not address these features in the final Office Action.** More particularly, the Examiner does not indicate how IVERSON et al. discloses or suggests a first policy using a first bucket to allocate the guaranteed bandwidth, a second policy using a second bucket to allocate the guaranteed

bandwidth, and borrowing bandwidth from the shared bandwidth bucket by one of the first and second buckets when the respective bucket has insufficient bandwidth to allow traffic to pass immediately through the network appliance. Rather, the Examiner merely indicates that “claim 14 is rejected for similar reasons as stated above.” (final Office Action -- pg. 3). The above identified features of claim 14 are not present in claim 1 and a rejection thereof is not supported based on a rejection of claim 1. Accordingly, a *prima facie* case of anticipation has not been established with respect to claim 14.

For at least the foregoing reasons, Appellants submit that the rejection of claim 14 under 35 U.S.C. § 102(e) based on IVERSON et al. is improper. Accordingly, Appellants request that the rejection be reversed.

B. The rejection under 35 U.S.C. § 103(a) based on IVERSON et al. (U.S. Patent No. 6,052,379) in view of HO (U.S. Patent No. 6,862,270 should be reversed.

The initial burden of establishing a *prima facie* basis to deny patentability to a claimed invention always rests upon the Examiner. *In re Oetiker*, 977 F.2d 1443, 24 U.S.P.Q.2d 1443 (Fed. Cir. 1992). In rejecting a claim under 35 U.S.C. § 103, the Examiner must provide a factual basis to support the conclusion of obviousness. *In re Warner*, 379 F.2d 1011, 154 U.S.P.Q. 173 (CCPA 1967). Based upon the objective evidence of record, the Examiner is required to make the factual inquiries mandated by *Graham v. John Deere Co.*, 86 S.Ct. 684, 383 U.S. 1, 148 U.S.P.Q. 459 (1966). *KSR International Co. v. Teleflex Inc.*, 550 U.S. ____ (April 30, 2007). The Examiner is also required to explain how and why one having ordinary skill in the art would have been realistically motivated to modify an applied reference and/or combine applied references

to arrive at the claimed invention. Uniroyal, Inc. v. Rudkin-Wiley Corp., 837 F.2d 1044, 5 U.S.P.Q.2d 1434 (Fed. Cir. 1988).

1. Claims 2, 3, 7-11, and 13.

Claims 2, 3, 7-11, and 13 depend from claim 1. The disclosure of HO does not cure the deficiencies in the disclosure of IVERSON et al. identified above with respect to claim 1. Therefore, claims 2, 3, 7-11, and 13 are patentable over IVERSON et al. and HO, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 1.

2. Claim 15.

Claim 15 is directed to an apparatus for allocating bandwidth in a network appliance where the network appliance includes a plurality of guaranteed bandwidth buckets used to evaluate when to pass traffic through the network appliance, the apparatus. The apparatus includes a shared bandwidth bucket associated with a plurality of the guaranteed bandwidth buckets; means for allocating bandwidth to the shared bandwidth bucket based on the underutilization of bandwidth in the plurality of guaranteed bandwidth buckets; and a scheduler operable to evaluate a packet to determine if a traffic shaping policy should be applied to a given packet, evaluate a guaranteed bandwidth bucket associated with an identified traffic shaping policy, determine when the guaranteed bandwidth bucket associated with an identified traffic shaping policy has insufficient capacity to support a transfer of the packet through the network, and borrow bandwidth from the shared bandwidth bucket by a respective guaranteed bandwidth bucket to allow traffic to pass immediately through the network

appliance. The combination of IVERSON et al. and HO do not disclose or suggest this combination of features.

For example, IVERSON et al. and HO do not disclose or suggest a shared bandwidth bucket associated with a plurality of guaranteed bandwidth buckets. Although not explicitly state in the final Office Action, the Examiner appears to again rely on the Abstract, Fig. 10, and col. 17, line 56 to col. 18, line 19 of IVERSON et al. for allegedly disclosing this feature (final Office Action -- pg. 6 ("As per claim 15, as closely interpreted by the Examiner, Iverson in combination with HO teach all that is similar above in claim 1 as application to claim 15"). Appellants respectfully disagree with the Examiner's interpretation of IVERSON et al.

This Abstract section of IVERSON et al. discloses a first allotted bandwidth bucket and a second overflow bucket. Depending on whether a water level in the first bucket is above or below zero, a priority value is derived from the level of either the first bucket or the second bucket. This section of IVERSON et al. does not disclose or suggest a shared bandwidth bucket associated with a **plurality of the guaranteed bandwidth buckets**, as recited in claim 15.

As described above, col. 17, line 56 – col. 18, line 19 of IVERSON et al. discloses a leaky bucket priority scheme, wherein excess bandwidth credits for a first committed bandwidth bucket 402 are added to a second excess bandwidth bucket 404. The excess bandwidth stored in bucket 404 is then used when the level of the first bucket 402 drops below zero (a midpoint in the bucket). This section of IVERSON et al. does not disclose or suggest a shared bandwidth bucket associated with a plurality of the guaranteed bandwidth buckets, as recited in claim 15. Even assuming *arguendo* that IVERSON et

al. disclose a shared bandwidth bucket (e.g., second bucket 404) associated with a single guaranteed bandwidth bucket (e.g., first bucket 402), a point that Appellants do not concede, this association is clearly a one-to-one association, resulting in bandwidth overages from bucket 402 being applied to bucket 404 for subsequent use when the level of bucket 402 drops below zero. Contrary to this disclosure, claim 15 recites a shared bandwidth bucket being associated with a plurality of the guaranteed bandwidth buckets. By associating multiple guaranteed bandwidth buckets with a shared bandwidth bucket, traffic resources may be more optimally distributed. The disclosure of HO does not cure the deficiencies in the disclosure of IVERSON et al. with respect to claim 15.

For at least the foregoing reasons, Appellants submit that the rejection of claim 15 under 35 U.S.C. § 103(a) based on IVERSON et al. and HO is improper. Accordingly, Appellants request that the rejection be reversed.

3. Claims 16-19.

Independent claim 16 is directed to a network device including a first bucket configured to receive tokens at a first information rate; a second bucket configured to receive tokens at a second information rate; a third bucket configured to receive extra tokens from the second bucket; and a scheduler configured to: determine if a size of traffic received at the network device exceeds a number of tokens stored in the first bucket, determine, when the size of the traffic does not exceed the number of tokens stored in the first bucket, if a size of the traffic exceeds a number of tokens stored in the second bucket, and transfer, when the size of the traffic exceeds the number of tokens stored in the second bucket, an appropriate number of tokens from the third bucket to the second bucket so that the second bucket includes a number of tokens that equals or

exceeds the size of the traffic. IVERSON et al. and HO do not disclose or suggest the combination of features recited in claim 16, either alone or in any reasonable combination.

For example, IVERSON et al. and HO do not disclose or reasonably suggest a first bucket configured to receive tokens at a first information rate; a second bucket configured to receive tokens at a second information rate; and a third bucket configured to receive extra tokens from the second bucket. The Examiner alleges that IVERSON et al.'s CIR 400 equates to the claimed first bucket, first bucket 402 equates to the claimed second bucket, and second bucket 404 equates to the claimed third bucket and relies on col. 17, line 41 to col. 18, line 20 of IVERSON et al. for allegedly disclosing these features (final Office Action – pg. 7. Appellants respectfully disagree with the Examiner's interpretation of IVERSON et al.

Col. 17, line 41 to col. 18, line 20 discloses

At the end of every evaluation interval the Committed Information Rate (CIR) quantum is emptied into a the CSum bucket 402 and/or the ESum bucket 404. The committed burst bandwidth credit (B_c) dimension of the first bucket 402 represents the amount of bandwidth that a User may transmit in a burst, potentially above the CIR, and expect reliable delivery to the network. The water level of the first bucket (BpCSum) represents the amount of bandwidth accumulated by the user above the CIR rate up to the maximum provisioned for the user (B_c).

Thus, if the BpCSum is stable from interval to interval, the User is requesting traffic delivery at their CIR. If the BpCSum rises from interval to interval, the User is requesting traffic at a rate below their CIR and if it is falling, the User is requesting traffic at a rate above their CIR.

If the BpCSum is positive, the port was requesting bandwidth at a rate below the CIR+ B_c for at least the last measurement interval. If the BpCSum is zero, port bandwidth requests have been substantially equal to the CIR+ B_c for the port. If the water level in CSum is negative (below the midpoint), the rate that the port has been using bandwidth is above

CIR+B_c. If the port has accumulated any excess bandwidth credit by transmitting below CIR for some amount of time, this bandwidth credit will be used if the water level in the first bucket goes below zero.

BpEsum is the water level value in the second bucket 404 and represents the current accumulated value of unused bandwidth in excess of CIR+B_c (i.e. past overflows from the first bucket 402). The ESum bucket 404 represents a cache of excess bandwidth that the user 62 can save up to be used for longer periods of high transmission demand.

Every measurement interval the quantum of bits 400 are added to the first bucket 402. Any overflow of bandwidth above the limit of the first bucket 402 is added to the ESum bucket 404.

Both buckets are "leaky" in that the amount of traffic transmitted in the past measurement interval leaks out of the appropriate bucket based on the previous priority level. The current water level of each bucket is then the result of adding in the Committed Information Rate (CIR) bit quantum for the last measurement interval and subtracting the amount of outgoing traffic 409 actually transmitted in the last measurement interval, T1Out. The water level of bucket 402 determines a priority value in a high priority band 403. The water level of bucket 404 determines a priority value in a low priority band 405.

This section of IVERSON et al. discloses a leaky bucket priority scheme, wherein excess bandwidth credits for a first committed bandwidth bucket 402 are added to a second excess bandwidth bucket 404. The excess bandwidth stored in bucket 404 is then used when the level of the first bucket 402 drops below zero (a midpoint in the bucket). The Committed Information Rate (CIR) is the information rate at which bits are assigned to the first bucket 402 for use by a user. The CIR is not a "bucket" that is assigned bandwidth at a first information rate. The second bucket 404 of IVERSON et al. is then assigned bandwidth left over from that assigned to bucket 402 at the CIR. Clearly, IVERSON et al. discloses only a single bucket (i.e., bucket 402) that receives bandwidth at a first information rate (CIR) and a shared bucket (i.e., bucket 404) that receives extra bandwidth from the first bucket. IVERSON et al. does not disclose or suggest a second

bucket that receives tokens or bandwidth at a second information rate and a third bucket that receives extra tokens or bandwidth from the second bucket.

Even assuming *arguendo* that IVERSON et al. discloses first, second, and third buckets (a point that Appellants do not concede), clearly IVERSON et al. does not disclose or even remotely suggest the first bucket receiving tokens or bandwidth at a **first information rate** and the second bucket that receives tokens or bandwidth at a **second information rate**, as recited in claim 16. The final Office Action is completely silent with respect to this feature. Accordingly, a *prima facie* case of obviousness has not been established with respect to claim 16. The disclosure of HO does not cure the deficiencies in the disclosure of IVERSON et al. with respect to claim 16.

Claims 17-19 depend from claim 16 and, as such, include each and every limitation included within the claims from which they depend. Therefore, Appellants submit that claims 17-19 are patentable over IVERSON et al. and HO, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 16.

Independent claim 20 recites features similar to (yet possibly different in scope than) features recited above with respect to claim 16. Therefore, Appellants submit that claim 20 is patentable over the combination of IVERSON et al. and HO, whether taken alone or in any reasonable combination, for reasons similar to reasons given above with respect to claim 16.

Claims 21 and 22 depend from claim 20 and, as such, include each and every limitation included within the claims from which they depend. Therefore, Appellants submit that claims 21 and 22 are patentable over IVERSON et al. and HO, whether taken

alone or in any reasonable combination, for at least the reasons given above with respect to claim 20.

For at least the foregoing reasons, Appellants submit that the rejection of claim 16-22 under 35 U.S.C. § 103(a) based on IVERSON et al. and HO is improper. Accordingly, Appellants request that the rejection be reversed.

C. The rejection under 35 U.S.C. § 103(a) based on IVERSON et al. (U.S. Patent No. 6,052,379) in view of Applicants' admitted prior art should be reversed.

1. Claim 4.

Claim 4 depends from claim 1. The disclosure of Appellants' allegedly admitted prior art does not remedy the deficiencies in the disclosure of IVERSON et al. set forth above with respect to claim 1. Therefore, Appellants submit that claim 4 is patentable over IVERSON et al. and Appellants' allegedly admitted prior art, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 1.

D. The rejection under 35 U.S.C. § 103(a) based on IVERSON et al. (U.S. Patent No. 6,052,379) in view of in view of CHIRUVOLU (U.S. Patent No. 6,839,321) should be reversed.

1. Claim 12.

Claim 12 depends from claim 1. The disclosure of CHIRUVOLU does not remedy the deficiencies in the disclosure of IVERSON et al. set forth above with respect to claim 1. Therefore, Appellants submit that claim 12 is patentable over IVERSON et

al. and CHIRUVOLU, whether taken alone or in any reasonable combination, for at least the reasons given above with respect to claim 1.

VIII. CONCLUSION

In view of the foregoing arguments, Appellants respectfully solicit the Honorable Board to reverse the Examiner's rejections of claims 1-22 under 35 U.S.C. §§ 102 and 103.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account No. 50-1070 and please credit any excess fees to such deposit account.

Respectfully submitted,

HARRITY SNYDER, L.L.P.

By: Robin C. Clark
Robin C. Clark
Registration No. 40,956

Date: May 9, 2007

11350 Random Hills Road
Suite 600
Fairfax, Virginia 22030
(571) 432-0800

Customer No. 44987

IX. CLAIM APPENDIX

1. A method for allocating bandwidth in a network appliance where the network appliance includes a plurality of guaranteed bandwidth buckets used to evaluate when to pass traffic through the network appliance, the method comprising:

providing a shared bandwidth bucket associated with each of the plurality of the guaranteed bandwidth buckets;

allocating bandwidth to the shared bandwidth bucket based on the underutilization of bandwidth in any one of the plurality of guaranteed bandwidth buckets;

determining whether bandwidth in one of the plurality of guaranteed bandwidth buckets is sufficient to allow traffic to pass immediately through the network appliance; and

transferring bandwidth from the shared bandwidth bucket to one of the plurality of guaranteed bandwidth buckets when it is determined that bandwidth in one of the plurality of guaranteed bandwidth buckets is not sufficient to allow traffic to pass immediately through the network appliance.

2. The method of claim 1 wherein the shared bandwidth bucket is a token bucket.

3. The method of claim 1 wherein the guaranteed bandwidth buckets are token buckets.

4. The method of claim 1 wherein the guaranteed bandwidth buckets are credit/debit buckets.
5. The method of claim 1 wherein each guaranteed bandwidth bucket is associated with a traffic shaping policy.
6. The method of claim 1 wherein a plurality of guaranteed bandwidth buckets are associated with a single traffic shaping policy.
7. The method of claim 5 wherein the traffic shaping policy screens based on IP address.
8. The method of claim 7 wherein the traffic shaping policy screens based on source IP address.
9. The method of claim 7 wherein the traffic shaping policy screens based on destination IP address.
10. The method of claim 7 wherein the traffic shaping policy screens based on protocol type.

11. The method of claim 7 wherein the traffic shaping policy screens based on UDP/TCP port number.

12. The method of claim 7 wherein the traffic shaping policy screens based on the type of service requested.

13. The method of claim 5 wherein the traffic shaping policy screens based on traffic content.

14. A method for allocating bandwidth in a network appliance comprising:
defining a guaranteed bandwidth allocation for a first policy for passing traffic through the network appliance including using a first bucket to allocate the guaranteed bandwidth;

defining a guaranteed bandwidth allocation for a second policy for passing traffic through the network appliance including using a second bucket to allocate the guaranteed bandwidth;

sharing excess bandwidth developed from the underutilization of the guaranteed bandwidth allocated to the first and second buckets including

providing a shared bandwidth bucket associated with the first and second buckets;
and

borrowing bandwidth from the shared bandwidth bucket by one of the first and second buckets when the respective bucket has insufficient bandwidth to allow traffic to pass immediately through the network appliance.

15. An apparatus for allocating bandwidth in a network appliance where the network appliance includes a plurality of guaranteed bandwidth buckets used to evaluate when to pass traffic through the network appliance, the apparatus comprising:

 a shared bandwidth bucket associated with a plurality of the guaranteed bandwidth buckets;

 means for allocating bandwidth to the shared bandwidth bucket based on the underutilization of bandwidth in the plurality of guaranteed bandwidth buckets; and

 a scheduler operable to

 evaluate a packet to determine if a traffic shaping policy should be applied to a given packet,

 evaluate a guaranteed bandwidth bucket associated with an identified traffic shaping policy,

 determine when the guaranteed bandwidth bucket associated with an identified traffic shaping policy has insufficient capacity to support a transfer of the packet through the network, and

 borrow bandwidth from the shared bandwidth bucket by a respective guaranteed bandwidth bucket to allow traffic to pass immediately through the network appliance.

16. A network device comprising:
 - a first bucket configured to receive tokens at a first information rate;
 - a second bucket configured to receive tokens at a second information rate;
 - a third bucket configured to receive extra tokens from the second bucket; and
 - a scheduler configured to:
 - determine if a size of traffic received at the network device exceeds a number of tokens stored in the first bucket,
 - determine, when the size of the traffic does not exceed the number of tokens stored in the first bucket, if a size of the traffic exceeds a number of tokens stored in the second bucket, and
 - transfer, when the size of the traffic exceeds the number of tokens stored in the second bucket, an appropriate number of tokens from the third bucket to the second bucket so that the second bucket includes a number of tokens that equals or exceeds the size of the traffic.

17. The network device of claim 16 wherein the scheduler is further configured to:
 - cause the traffic to be forwarded after the transfer; and
 - decrement the number of tokens in the first and second buckets based on the size of the traffic.

18. The network device of claim 16 wherein the scheduler is further configured to:

determine if the third bucket includes the appropriate number of tokens, and prohibit the traffic from being forwarded when the third bucket includes less than the appropriate number of tokens.

19. The network device of claim 16 further comprising:
one or more input ports configured to receive traffic from a network, each of the one or more input ports including the first bucket, the second bucket, the third bucket, and the scheduler.

20. A method comprising:
receiving traffic;
determining if a policy is to be applied to the traffic;
determining, when a policy is to be applied to the traffic, if a size of the traffic exceeds a number of tokens in a first bucket, the first bucket being associated with the policy;
determining, when the size of the traffic does not exceed the number of tokens in the first bucket, if the size of the traffic exceeds the number of tokens in a second bucket;
determining, when the size of the traffic exceeds the number of tokens in the second bucket, if a third bucket includes an appropriate number of tokens that, when

added to the number of tokens in the second bucket, would equal or exceed the size of the traffic;

transferring the appropriate number of tokens from the third bucket to the second bucket when the third bucket includes the appropriate number of tokens; and

forwarding the traffic after the transferring.

21. The method of claim 20 further comprising:

forwarding the traffic when the size of the traffic does not exceed the number of tokens in the second bucket.

22. The method of claim 20 further comprising:

repeating the determining if a size of the traffic exceeds a number of tokens in a first bucket; the determining, when the size of the traffic does not exceed the number of tokens in the first bucket, if the size of the traffic exceeds the number of tokens in the second bucket;

the determining, when the size of the traffic exceeds the number of tokens in the second bucket, if a third bucket includes an appropriate number of tokens that, when added to the number of tokens in the second bucket, would equal or exceed the size of the traffic; and the transferring the appropriate number of tokens from the second bucket to the first bucket when the third bucket includes the appropriate number of tokens for at least a second policy prior to transferring the traffic, the second policy being associated with different first and second buckets.

APPEAL BRIEF

PATENT

U.S. Application No. 09/658,424

Attorney's Docket No. **0023-0200**

X. EVIDENCE APPENDIX

None.

XI. RELATED PROCEEDINGS APPENDIX

None.